

EXECUTIVE OFFICE OF THE PRESIDENT
OFFICE OF SCIENCE AND TECHNOLOGY POLICY
WASHINGTON, D.C. 20500

January 25, 1983

MEMORANDUM FOR MEMBERS FCCSET SUPERCOMPUTER PANEL

FROM: DOUG PEWITT *Doug*
SUBJECT: Follow-up to January 24, 1983 Panel Meeting

The following are the three questions we agreed to consider and to provide our individual agency views on for circulation by the week of January 31:

- i) What the Federal Government should do, if anything, to meet the declared intent of the Japanese initiatives to attain world leadership in the supercomputer field, and to meet our own needs for very high speed, general purpose supercomputers.
- ii) How to insure access to existing supercomputers for users whose present research computing needs cannot be accommodated on less than Class VI machines.
- iii) How to coordinate Federal R&D activities in artificial intelligence, Fifth Generation computers, and information processing related technologies generally.

We agreed that the last two questions might be appropriately left to technical working groups, but we need to agree on whether they are necessary, and if so, on our charges to them.

At our next meeting we will discuss proposed solutions to the first question. Hopefully, we can agree on what the Administration's response to this challenge will be, even if no common strategy is advisable for all agencies.

Attached is a listing of attendees with addresses.

LIST OF ATTENDEES

COLLADAY, Raymond
Deputy Associate Administrator
National Aeronautics and Space Administration
Washington, D.C. 20546

COOPER, Robert
Director
Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, Virginia 22209

STAT

Director
Office of Planning
Intelligence Community Staff
Washington, D.C. 20505

EDEN, Frank
Senior Science Associate
National Science Foundation
1800 G Street, N.W.
Washington, D.C. 20550

KAHN, Robert
Director
Information Processing Techniques Office
Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, Virginia 22209

KENEBROCK, Jack
Associate Administrator
for Aeronautics and Space Technology
Code R
National Aeronautics and Space Administration
Washington, D.C. 20546

PEWITT, N. D.
Assistant Director
Office of Science and Technology Policy
The White House
Washington, D.C. 20500

SMAGORINSKY, Joseph
Geophysical Fluid Dynamics Laboratory
National Oceanic and Atmospheric Administration
Post Office Box 308
Princeton, New Jersey 08540

TRIVELPIECE, Alvin
Director of Energy Research
Room 7B-058
Department of Energy
Forrestal Building
Washington, D.C. 20585

WELSH, James
Supervisory Computer Systems Analyst
Geophysical Fluid Dynamics Laboratory
Princeton University
Post Office Box 308
Princeton, New Jersey 08540


PETTIFORE, Andrew
Policy Analyst
Office of Science and Technology Policy
The White House
Washington, D.C. 20500

**EXECUTIVE OFFICE OF THE PRESIDENT
OFFICE OF SCIENCE AND TECHNOLOGY POLICY**

WASHINGTON, D.C. 20500

March 18, 1983

MEMORANDUM FOR DISTRIBUTION


FROM: DOUG PEWITT 
SUBJECT: FCCSET Supercomputer Panel Feedback

Attached are copies of the responses you have each provided so far, as agreed to at our last meeting.

It seems clear that the government has a role both as a customer for the latest "off the shelf" systems, and as a friendly host for supercomputer prototype evaluations. In addition, the government has a major role in supporting, at a minimum, supercomputer relevant, non proprietary research and the associated training of future computer scientists.

Our next meeting has been scheduled for 3:00 p.m., March 30, in Room 5026. At that meeting, I would like to discuss how best we should implement these or other appropriate roles. This would enable us to reach a conclusion, on a plan of action at least, to the first of the three questions we posed at the last meeting. If there is time we will discuss how best we should approach the subsequent questions of user access and research coordination.

Attachments

R. Cooper, DOD /
G. 
F. Eden, NSF
J. Kerrebrock, NASA
J. Smagorinsky, NOAA
A. Trivelpiece, DOE

STAT



Department of Energy
Washington, D.C. 20585

MAR 8 1983

MEMORANDUM FOR: N. Douglas Pewitt

SUBJECT: Department of Energy Views on Supercomputers

Background

The Department of Energy's (DOE) laboratories have historically required the use of the fastest computers available in order to carry out their respective missions. Since the 1950's, a DOE laboratory has acquired the first or one of the first manufactured units of nearly every new large scale computer. These early models arrived almost devoid of usable software and consequently the laboratories were required to provide the necessary systems and application programs in order to fully utilize the new machines. Thus, DOE laboratories have played an important role in causing the development of each new generation of supercomputers.

Every generation of large scientific computers has had a significant impact on the ability of the DOE laboratories to meet their mission goals. Each new machine has improved the laboratories' abilities to more accurately model complex physical phenomena thus improving their predictive capability and reducing dependence on empirical scaling. This is particularly true in the design and testing of nuclear weapons and more recently in controlled nuclear fusion research. A recent inventory * of supercomputers (Cray 1 and Cyber 205) indicates that 36% of the total number in the U.S. are located in DOE laboratories.

Until the late 1960's, a new computer generally provided an order of magnitude (10x) increase in capability. Since that time, however, the increase has been more like fourfold (4x) with each new major main frame. At the same time, the computational needs of the DOE laboratories have increased at a faster rate. Consequently, supercomputer technology has fallen short of meeting present desired capability. Today, DOE laboratories have major programs that have a strong need for a minimum of a 200x increase in capability by the end of the decade. Current hardware development schedules by U.S. manufacturers further indicate that far less than this improvement will be achieved by 1989 unless a significant major increase in development effort occurs in this country.

*Report of the Panel on Large Scale Computing and Engineering,
 Peter D. Lax (Chairman), NSF, December 1982.

Although scientists and engineers in DOE's weapons and magnetic fusion programs have ready access to state-of-the-art (Class VI) supercomputers, scientists in other DOE programs such as Basic Energy Sciences, High Energy Physics, Nuclear Physics and others do not have such access. A similar situation exists in research and development programs supported by other U.S. government funding agencies. Recent studies * suggest that many of these R&D programs would benefit substantially from the use of supercomputers. Lack of access by universities causes a special problem since there is now very little training of new scientists and engineers in the use of supercomputers. In DOE, the Office of Energy Research is presently considering alternatives for providing supercomputer access for its scientists. As an interim solution 800 hours/per year of Cray 1 time on the Magnetic Fusion Energy Network has been made available to other scientists in the Department's Energy Research programs in DOE laboratories and universities.

Both the application of supercomputers to present R&D programs and the development of new generations of supercomputers are important to the future leadership of the U.S. in defense technology; in research and development; and in many commercial areas as well. Perhaps one of the most important issues is the national security implication of buying supercomputers from a foreign country if U.S. industry were to lose its leadership in this critical area. For these reasons, we believe that the U.S. government should take appropriate steps to encourage application of present supercomputers to a variety of scientific and engineering problems and to stimulate U.S. computer vendors to accelerate development of next generation machines.

Recommended Actions

1. Accelerated Development of New Supercomputers.

Should the U.S. government play a role in the development of supercomputers? There are several facts about the supercomputer industry that are important to this discussion. The present supercomputer market is relatively small at about \$100M per year. It appears that most of the large computer manufacturers have decided not to enter this market segment presumably because of its small size. Only two companies, CDC and Cray Research, have produced Class VI machines.

Start-up and R&D costs could range from \$50-100M for a new generation of supercomputers. Cray Research is a relatively small company and Denelcor, another vendor which could potentially become a force in the supercomputer market is a very small company. The R&D capabilities of both Cray and Denelcor are limited and it is questionable whether they can compete with government subsidized vendors in other countries in the long term. In summary, the large computer companies seem to be reluctant to devote substantial R&D funds to supercomputers and the smaller companies have very limited resources for these developments.

*Report of the Panel on Large Scale Computing and Engineering,
Peter D. Lax (Chairman), NSF, December 1982.

Another important fact is that the U.S. government is a large consumer of supercomputers having purchased or leased over 50% of all of the Class VI computers currently in service in the United States. The government has played a significant role in developing useful supercomputer systems in the past to meet its own needs. For example, the DOE accepted the risks by buying the first production Class VI machines and the DOE laboratories did substantial software development as they have historically. Other customers started purchasing Class VI machines only after they had been demonstrated in DOE laboratories.

In view of the U.S. government's need for advanced supercomputers, the limited development capabilities of some key vendors, and the historical government role in development of supercomputers, several government actions are recommended. If the government plays the proper role, it could stimulate competition within the United States as well as ensuring the continuing dominance of the United States in future supercomputer markets.

A. Procurement Actions

Both Cray Research and Control Data Corporation are developing advanced machines. The first of the new Cray-2 machines is expected to be available in early CY 1985. The CDC machine is expected sometime later. For the near term, the government should continue to play the role of purchasing early production models and its laboratories should continue to assist in the development of software for those production models.

For the generation of machines beyond the Cray-2, the government should provide incentives that will accelerate supercomputer development in order to meet its own needs and to help ensure continued U.S. leadership in this critical field. A goal of at least 200x present capability by 1989 appears feasible, but attaining this without government incentives is unlikely. The target machines would be commercially viable, with a selling price of \$10-15M. The Japanese have formally created a national program to achieve a supercomputer with 100x present capability by 1990.

The program proposed for the U.S. government is based upon a phased, competitive procurement. The details of this procurement have not been worked out, but it might have three phases as described below.

The first phase would solicit proposals to generate paper specifications from a number of potential vendors. This phase would last one year and might involve 5-6 vendors. The product of this phase would be a number of paper designs of supercomputer systems, including architecture, technology, and software. Evaluation of this phase would benefit from the experience gained with experimental machines, discussed in IB.

During the second phase, the most promising designs (perhaps three) would be supported for prototype hardware and software development. The second phase would last two years and the government would receive the right to use the prototype.

During phase three, those prototypes which were judged capable of attaining the performance goals would undergo final development and initial production. This phase, lasting two years, would also include experiments using the prototype machines by interested government laboratories in close cooperation with the vendors. Government funding of phase three would be through a guaranteed purchase of approximately two machines per vendor with the expectation that there would be additional purchases by both government and industry. Through this mechanism both the hardware and software may be modified to better suit the end user's need.

At the end of the approximately five year development cycle, the new machine(s) becomes fully commercial. If this process is successful, it is envisioned to lead to a semi-continuous process that could be repeated for the following generations of supercomputers.

B. Government Supported Research and Development

In addition to the procurement activities described, which will stimulate industrial R&D, some specific government supported R&D is recommended. The government should selectively provide early financial support to innovative architectural and technological approaches taken by private industry. This would help to ensure survival of such approaches until they can become marketable as production solutions to known requirements. It is also needed to encourage private industry to engage in high risk R&D ventures which may accelerate the advancement of supercomputers.

One approach the government should consider is the immediate establishment of an R&D program(s) to purchase experimental systems such as the Denelcor Heterogeneous Element Processor (HEP) and the CDC Advanced Flexible Processor (AFP). Several of these experimental systems should be purchased and installed in various federal laboratories and made available for research purposes to universities and industry as well as the laboratories. This would permit research on the use of advanced architectures as they relate to new languages, algorithms, software tools and applications systems.

Care should be taken to assure that a clear distinction is made between the acquisition and use of such experimental systems versus the acquisition and use of production computers to avoid misunderstandings concerning the expectation of results.

Future increases of orders of magnitude in supercomputer speeds can be achieved only by incorporating multiple processors and perhaps hundreds or thousands of processors in new parallel architectures will be required. For approximately the next five years, the appropriate role for DOE is to support long term research in languages, algorithms, and software for parallel architectures. Government agencies should provide support for: faculty, staff, and graduate students; access to commercially available experimental systems; and several university/industry/laboratory collaborations to build

and experiment with new experimental hardware/software systems. Increased support for some of these longer-range research activities will be required if they are to impact the development of commercial machines in a timely fashion.

Since the DOE is the largest user of supercomputers and has traditionally led in advancing the state-of-the-art in numerical algorithms and software, the DOE should remain the lead government agency in accelerating the development of advanced supercomputers and in exploration of large scale scientific computational techniques.

II. Providing access to supercomputers for a larger number of scientists and engineers.


Providing access to present and future supercomputers for a larger number of U.S. scientists and engineers is an important step that will provide several benefits. Supercomputers are important new scientific tools which have the potential to advance rapidly and substantially research and development in many fields. Making supercomputers available to researchers in universities has the additional benefit that students will become familiar with their use. The training aspect is important because presently very few scientists and engineers graduate with any experience in using this modern scientific tool. Also, greater use of supercomputers will undoubtedly increase their applications and should help bring solutions to many problems where their use may not now be anticipated. These increased applications would expand the supercomputer market.

In DOE's magnetic fusion program, which does provide access to Class VI computers for all its contractors, substantial benefits have been derived. Other programs and other agencies should carefully consider the potential benefits of greater access to existing supercomputers for researchers in their R&D programs. This should be an individual agency decision. As mentioned above, DOE's Office of Energy Research is currently planning to provide access for more of its scientists in DOE laboratories and universities.

A FCCSET working group should be formed to discuss potential interagency cooperation and to share experiences and plans in this area.

III. Coordination of Federal R&D Activities in Artificial Intelligence, Fifth Generation Computers and Information Processing.

There is currently research and development in advanced computer architecture, hardware, software, etc., funded by several government agencies. It would undoubtedly be beneficial to coordinate this R&D across the government. An interagency committee of working level people should be established to coordinate activities in advanced computer development.


Alvin W. Trivelpiece
Director, Office of
Energy Research



Dr. N. D. Pewitt
Assitant Director, OSTP
The White House
Washington, DC 20500

U.S. DEPARTMENT OF COMMERCE
National Oceanic and Atmospheric Administration
ENVIRONMENTAL RESEARCH LABORATORIES
c/o GEOPHYSICAL FLUID DYNAMICS LABORATORY
Princeton University - P.O. Box 308
Princeton, New Jersey 08540

Phone: 609-452-6502, FTS and Commercial

February 9, 1983

R/E/GF

TO: Doug Pewitt
FROM: Joseph Smagorinsky (Home Phone: 609-921-2762)
SUBJECT: NOAA Recommendation for Supercomputer Policy

In my opinion the first of the three questions enumerated in your January 25, 1983 memo can be separated from the last two, which can be dealt with at a later time. I feel that question (ii) is important, but I do not presume to contribute much to question (iii).

My views on Japanese competition are briefly addressed in section A, below, in the form of questions that it raises on motivation and ultimate objectives. The remaining sections discuss NOAA's needs and procurement practices as they may be relevant to an evolving Government-wide policy. They are proposed for consideration to stimulate the best and most timely response from the industry, whether the market place be world-wide or confined to American vendors.

A. U.S. Industry vs. World Competition. There is little question that the U.S. Government is the major user of top-of-the-line computing systems and will continue to be so. It will therefore need from the industry at-large the maximum rate of improvement in power at the least possible cost. To this end the U.S. Government should assure a healthy competition in top-of-the-line computers.

Whether the U.S. Government seeks such a yield exclusively from American industry depends on several factors:

* Will it inhibit constructive competition in the form of rapid innovation and minimum price?

* Will it insure self sufficiency (in a national security sense) under a less friendly political situation?

* Is it useful to distinguish an open market for components only (e.g., I.C. chips) from an open market for integrated main frames?



Doug Pewitt; February 9, 1983; Page 2

* Should Government preference for American vendors be in the form of protective tariffs on Japanese imports or just in the form of a subsidy for American industry to stimulate development?

* If the Government subsidizes American industry, should it be through a direct subsidy (the way the Japanese do it) or indirectly by a guarantee to acquire a minimum number of early copies of each generation (the way AEC did it in the 1950's, 60's and 70's).

B. The Scope of Meteorological Needs. The use of supercomputers in meteorology historically goes back to some of the first research studies on the computer built at the Institute for Advanced Study (IAS) in Princeton in the early 1950's. One of the predecessors of NOAA, the U.S. Weather Bureau, jointly with the Air Force and Navy, soon after acquired an IBM 701 for operational short range numerical weather prediction to mark the beginning of an unbroken sequence of top-of-the-line computer acquisitions. By next year NOAA will have three CYBER 205's, dedicated to meteorological and oceanographic research and to operational weather prediction.

This demand by the geophysical fluid dynamics sciences is further reflected in the USA by the National Center for Atmospheric Research (Cray-1), which is NSF supported, by partial use at the NASA Goddard Space Flight Center (CYBER 205) and by the DOD (CYBER 203, upgraded to a 205). A few universities also serve academic meteorological research with supercomputers, e.g., the Colorado State University (CYBER 205). Abroad the U.K. Meteorological Office uses a CYBER 205, and the Canadian Meteorological Service and the European Center for Medium-Range Forecasts each has a CRAY-1. The Japan Meteorological Agency now has at least two top-of-the-line HITAC-M200H computers. One of these is scheduled to be replaced in the near future by a more advanced model. The Meteorological Department at Tokyo University has access to the University's HITAC computer.

In short, the meteorological research and operations community has become a major market for supercomputers both domestically and world-wide. Most calculations to integrate non-linear systems of three space-dimensional, time-dependent systems of partial differential equations are still severely resolution limited. Since just a doubling of resolution in each of the three space dimensions generally entails a 16-fold increase in the amount of computation, it is likely that such uses will continue to be at the forefront of demand for supercomputers of the future. This, of course, is a direct counterpart to the needs of "computer wind tunnel"

Doug Pewitt; February 9, 1983; Page 3

applications. The main difference is that the number of systems needed in meteorology and oceanography may be an order of magnitude greater.

C. Software. There is fundamental significance in the use of the term "system", in that software capability can influence throughput by as much as an order of magnitude. The term "supercomputer" is not well focused. It seems to draw more attention to raw hardware speed than to programmability. It is therefore important to include the significance of software in the operative description, for almost all current supercomputer applications are, in fact, programmed in Fortran and run on more-or-less conventional operating software. The computer system may well be termed "Fortran engine". Software specifications to bridge the gap between hardware improvement and enhanced user needs are an integral element of system upgrading, whether the software is provided by the hardware vendor, by another commercial vendor or by the user itself.

D. Industry Responsiveness: Off the Shelf vs. Subsidized Development. NOAA has historically favored the normal commercial development of supercomputers. The most effective policy that can be followed by the U.S. Government is to assure a stable marketplace for such machines, although a stable, better-defined marketplace will not guarantee continued U.S. strength. The principal reason that the current supercomputer marketplace is not as effective as possible is the great inefficiency of the Federal procurement process and the balkanization of requirements. Each procurement of a supercomputer is a separate, isolated process that takes typically 4-5 years from requirements to delivery. Because each Federal agency must completely perform its own procurement, its specifications tend to be particular and self-serving, which thwarts standardization and cannot benefit from the clout of potential multiple procurements.

NOAA would be well served by a Federal policy that calls for annually competed supply contracts for top-of-the-line computing systems. By eliminating the RFP-Proposal cycle, agencies should be able to acquire such systems in 2 years, instead of 4-5. A top-of-the-line computing system is not expected to stand alone; a similar procurement mechanism should be in place for a time-sharing system that can serve as a front end to such computing engines.

A government-wide supply contract for top-of-the-line computing systems would have the following advantages:

Doug Pewitt; February 9, 1983; Page 4

* By reducing the procurement cycle significantly, the Government will benefit more rapidly from technological advances and greater efficiency available in the marketplace.

* Hardware and software standards can be more easily enforced both on vendors and on the agencies.

* Vendors can anticipate a predictable, steady market.

* Vendors will have as a design target a standard functional architecture that is based on Fortran.

E. Maximum Bang for Buck vs. Minimum Price. It is virtually axiomatic that a top-of-the-line customer wants the most powerful system available at the time. It is of interest and significance to note that comparison of today's supercomputer (say the CYBER 205, or CRAY-1) with that of 30 years ago (the IBM 701) shows less than a factor of three increase in inflation-adjusted dollars. For this, the present generation is 10,000 times more powerful, has 1000 times more high-speed memory, has almost limitless disk storage, and possesses an extremely sophisticated software ensemble which is about half the cost of the system.

It is proposed that supply contracts for state-of-the-art computer systems be awarded as the result of annual competitions. Agencies would then be free to place orders subject only to the availability of funds. For the agency the entire RFP-proposal-award cycle is bypassed. Delivery should be possible within 12 months after receipt of order.

The basic idea for evaluating competing systems is to select the system that can demonstrate the most productivity on a standard benchmark suite. The method of acquisition for evaluation should be lease. One of the goals of the streamlined procurement process is to facilitate more frequent turnaround of systems, so that the Government is always exploiting the latest technology and efficiency as much as possible. Whereas the conventional procurement fixes workload and competes on lowest price; the strategy for computer systems should fix price and compete on highest productivity. Fixing the price of supercomputers facilitates long-term budgetary and programmatic planning throughout the Government.

NOAA's recent procurement that resulted in a system with two CYBER 205's was conducted on this basis of most productivity for a specified annual budget. The only problems we have encountered are associated with relaxing

Doug Pewitt; February 9, 1983; Page 5

standards and requirements to assure a basic level of competition (two vendors) for the procurement. In any government-wide competition any compromise of standards must be strictly avoided. It is precisely through the standards set for competition that the Government can most effectively influence the future of supercomputers.

F. Shared Risk. A fundamental problem with top-of-the-line computer systems is that the meeting of design specifications is unknown until after delivery. Such systems are often not available for actual testing at the time that the contract is let. A competitive choice is normally based on indirect preliminary inferences and claims by the vendors. The Government must be protected against irresponsible or inaccurate initial claims by not accepting the complete risk for non-performance. Contractual provision must be made for adjustments to be levied upon the vendor for not meeting specification claims. NOAA has successfully imposed such contractual constraints on two procurements during the past decade at GFDL: a Texas Instruments ASC/4X delivered in 1974 and a CYBER 205 system delivered in 1982. The T.I. contract provisions were actually enforced for non-performance and the CYBER procurement is now similarly being tested.

cc: Dr. A. Calio, Deputy Administrator, NOAA



National Aeronautics and
Space Administration

Washington, D.C.
20546

Reply to Attn of:

RT

FEB 3 1983

Mr. Doug N. Pewitt
Assistant Director
Office of Science and Technology Policy
The White House
Washington, DC 20500

Subject: NASA (OAST) Position on Government's Role In
Support of General Purpose Supercomputers

This is our response to point one of your inquiry of January 25, 1983.

We believe the government should have three distinguishable roles in furthering the development and utilization of supercomputers in the U.S.: (1) encouraging private development of advanced supercomputers by acting as a supportive first user, (2) supporting research in computer science and technology, and (3) assuring U.S. manufacturers free competitive access to international markets.

The first two roles are currently recognized by government agencies, however increased emphasis is warranted in light of the current technical opportunities and international competitive situations. The third role should be reexamined, in the context of overall U.S. trade policy.

Encouraging Development of Advanced Supercomputers

By acting as first users of prototype supercomputers, federal R&D agencies can assure early evaluation of technical opportunities and assist in the early recovery of private R&D investments. The prototype hardware and software will enhance the user agencies' capabilities in R&D and its use as early as possible will stimulate beneficial interaction between the user community and the developer. Enhanced capability to address many technical problems of critical importance will result. Since development

costs for prototype hardware and systems software are low relative to those in many other high technology industries, rapid recovery of a portion of development cost through early sales to government R&D agencies should stimulate the rate at which new computers can be introduced. To assure ready acceptance of new products by commercial users, it is essential that the government work with the manufacturers to mature prototype designs into commercial systems.

The government should encourage applications of supercomputers by supporting research leading to new algorithms and numerical methods for the efficient solution of large engineering and scientific problems. Such research, motivated by fluid mechanics, structures and dynamics, chemistry, and physics, can stimulate existing markets for supercomputers, and develop new ones. Advances in computational fluid dynamics can lead to the use of the supercomputer in the design and analysis process for complete aircraft; advances in computational chemistry could lead to the use of the supercomputer in the chemical industry to develop synthesis processes and to select catalysts to control reaction rates.

The government should increase its utilization of supercomputers to meet unique government needs in areas such as weather prediction, defense system analysis, nuclear energy, and space imaging and data analysis. Such applications can provide a market for increased numbers of supercomputers, while fostering the aims of the user agencies.

Finally, the government should provide low cost supercomputer time to critical industrial and academic users, who otherwise would be unable to afford it. Focal points within the government should not only provide access, but also provide supporting expertise so that new users can efficiently utilize the computer in the solution of their problems. Such use will not only encourage additional applications, but also provide feedback to the computer system developers of potential computer hardware and software advances which will enhance capabilities.

Computer Science Research

The government must take a strong role in support of academic, government, and industrial research in computer science and technology. The thrust should be to accelerate the solution of technical problems associated with components, system architectures, algorithms and software for advanced supercomputers. Research should range from basic phenomena through proof-of-concept prototypes.

Remove Constraints to International Trade

Competition with foreign producers of supercomputers is adversely affected by restricted access to foreign markets by U.S. super-computer manufacturers. Industry contends that an expanded market provides R&D finance essential to assure their ability to compete with foreign manufacturers. There is a compromise to be made between restricting the export of production supercomputers to limit foreign nations' exploration of their capabilities and improving the economic health of the domestic industry to help preserve its leading position in technology. Foreign sales of U.S. production supercomputers are sometimes restricted for strategic reasons by the U.S. government, and the voids will soon be filled by the development of supercomputers of foreign design and manufacture. These developments are fostered by foreign government investment with their computer industry. The government should reassess its policy toward trade restrictions to consider the value of maintaining a technological lead in commercial supercomputers by assuring a larger foreign market versus restricting sales to a smaller member of strategically safe countries and allowing entry of foreign developed super-computers into the market place.



Jack L. Kerrebrock
Associate Administrator for
Aeronautics & Space Technology



DEFENSE ADVANCED RESEARCH PROJECTS AGENCY

1400 WILSON BOULEVARD
ARLINGTON VIRGINIA 22209

JAN 31 1983

MEMORANDUM FOR DOUG PEWITT

SUBJECT: FCCSET Supercomputer Panel Comments

Doug -

The comments below reflect the views of the Defense Advanced Research Projects Agency on the three questions addressed in your memorandum of 25 January. I believe they also reflect, generally, the views of Defense but, in view of the short suspense, we have not had time to coordinate them.

- i) The government should continue to provide support to agencies such as DOE and NASA which are heavy users of general purpose supercomputers so they may buy them from industry at a rate which is sufficient to meet their needs and to retain a U.S. capability in the field. Except for specialized NSA requirements, Defense has not been a major user of conventional supercomputer technology. Defense anticipates requirements for a new kind of very fast symbolic computing capability which relies heavily on artificial intelligence and for which significant research is still needed. Industry does not currently have the know-how to build machines of this type.
- ii) The existing class VI machines may not have sufficient capacity to support all the users whose present research cannot be accommodated with less computing power. A survey should be conducted to determine the major class VI requirements within government, industry and university groups carrying out state-of-the-art scientific research. It may be appropriate to authorize the purchase and operation of additional machines over a suitable time frame for shared use by the community. These machines could be made more widely available by connection to an appropriate network subject to appropriate administrative controls on access to the machines. The National Laboratories at Livermore and Los Alamos may be the appropriate place to house this capability.
- iii) At the present time, there are few if any major development programs in government or industry in artificial intelligence or 5th Generation computers. If, as expected, R&D in this area should increase significantly over the next ten years, existing coordination groups such as the Interagency Committee on Extramural Mathematic Programs (ICEMAP) would be candidates to include it. NSF and DoD

officials also meet regularly to discuss R&D programs in both organizations and this would be another forum for increased coordination. No further coordination action appears warranted in this area at the present time.

A handwritten signature in black ink, appearing to read "B. Cooper".

Robert S. Cooper
Director



DEFENSE ADVANCED RESEARCH PROJECTS AGENCY
1400 WILSON BOULEVARD
ARLINGTON VIRGINIA 22209

FEB - 1 1983

Dr. N. D. Pewitt
Assistant Director
Office of Science and Technology Policy
The White House
Washington, D.C. 20500

Dear Doug:

This is a follow-up to your request for comment on the three questions on Supercomputers which you enunciated at the meeting last Monday. Although Bob Kahn prepared an "ARPA" response (which you will receive under separate cover), I thought I would take the opportunity to expand on these important questions personally.

From a Defense (historical) perspective the market for Class VI and beyond number crunchers is very limited. As I pointed out at Monday's meeting, Defense has not taken advantage of such machines historically except in an indirect way through the DOE weapon design community and through the intelligence agencies. Instead the vast majority of Defense fixed base computer acquisitions have been for less capable machines that have been used for data base manipulation and C³ functions.

Although the scientific and engineering uses of Class VI machines would benefit the Defense Research community, defense research managers have not felt that they could justify the levels of funding (\$10's of millions) in their budget advocacy process to buy supercomputers. Hence, they have made no serious attempts to acquire them. What is more, the barriers thrown in the way of computer acquisition by Federal Regulations, GSA and Congressman Brooks have had a chilling effect on all scientific computer acquisition overall.

A factor in the disinterest shown specifically by the Defense labs in Class VI computing machines is the lack of qualified personnel in the labs to use such machines. Government personnel rules including the ones regulating general schedule employees and the senior executive service employees salaries, perks and working conditions have driven many of the best persons in the computer science and engineering fields out of government. The labs no longer have extensive capability to provide high quality scientific talent and computer code for machines such as the CRAY-1 which is delivered with very rudimentary software systems.

The net conclusion I draw from all of this is that the Japanese could sieze the entire market for supercomputers and it would affect the Defense Research community not at all. Furthermore, in order to create an interest in supercomputers and a highly qualified user community in Defense Research, we would have to solve some nearly intractable problems of regressive Government personnel policy, restrictive federal regulation of computer acquisition and unbridled turfsmanship by GSA and Congress. These are probably worthy problems to be attacked by the "Man of LaMancha."

My view on your second question is that we have not provided appropriate computing resources across the board for the university scientific community including general purpose supercomputers. Both NSF and Defense have plans over a several year period of time to upgrade the computing machines generally available for basic research in the computer community. However, there are no general plans to acquire and make available Class VI and beyond machines. They are just too expensive and hence do not relate well to our basic research programs. Only a few specialized cases such as that at NCAR and at Ames Research Center have actual plans related to G.P. Supercomputers being developed.

I feel that the Nation would benefit from policy level support for supplying computer cycles to basic research institutions. It could help NSF and DARPA get both Administration support (OMB, SecDef, etc.) and congressional support for our programs. In addition, I believe a specific recommendation to make the CRAY machines at LASL and L³ available over a network to qualified researchers is in order. Both of the computer departments at these institutions are favorable to such an arrangement. If a survey reveals that legitimate demand for cycles is greater than supply, then DOE's budget should be augmented to acquire more resources to meet the demand.

On the question of so-called fifth generation computers and artificial intelligence research, it is probably too early to augment the existing coordination mechanisms on programs in this area. Little real program activity is planned in this area prior to FY84 and even then the ideas upon which new programs will be based will not be clear until FY85. It appears that the right strategy here is to let plans develop under existing mechanisms for about a year or so then take another look to see if something more is needed.

I hope these views are useful in your consideration of actions needed.

Sincerely,



Robert S. Cooper
Director

NATIONAL SCIENCE FOUNDATION

WASHINGTON, D.C. 20550



OFFICE OF THE
ASSISTANT DIRECTOR
FOR ASTRONOMICAL,
ATMOSPHERIC, EARTH,
AND OCEAN SCIENCES

February 7, 1983

MEMORANDUM

TO: D. N. Pewitt
Assistant Director, OSTP

FROM: H. Frank Eden *H. F. Eden*
NSF Member FCCSET Supercomputer Panel

SUBJECT: NSF Views on Supercomputer Topics

This memo outlines the NSF view of the three topics raised at the panel meeting of January 24 and in your memo of January 25.

Background: NSF Involvement

NSF is involved in large scale computing in two principal ways:

- Research in computer science and engineering is supported through the Divisions of Mathematical and Computer Sciences, Electrical, Computer and Systems Engineering and Information Science and Technology. Topics supported cover the entire spectrum of basic research in computer science and computer engineering.
- Through the National Center for Atmospheric Research (NCAR), NSF has made available to atmospheric and ocean scientists a computing facility based on the most advanced machines available over the past 20 years. Until 1983 the facility used a CDC 7600 and a CRAY I. Recently NCAR acquired a second CRAY I as a replacement for the CDC 7600 and the facility is converting to use both CRAY machines. The facility services some 900 visitors and 75 remote terminals across the Nation. Plans call for the acquisition of a Class VII machine to replace one of the CRAY I machines over the next few years.

- 2 -

- In addition, researchers in disciplinary areas requiring large scale computation can include computing costs in their proposals for NSF support as appropriate. In all disciplines NSF supports the acquisition of computing equipment.

The Federal Government and Future Supercomputers

The recent history of the industry is one of progression in the development of increasingly powerful general purpose machines. Class VII machines are already under development with, for example, CRAY II planned to be available in 1984-1985. It seems likely that this trend will continue and Class VIII machines will appear in the early 1990's unless likely vendors are discouraged from making the necessary investment in R and D. There is a limited market worldwide and stiff competition expected from Japan. Decisions by individual companies will be made on the basis of their internal assessment of the likely market which will be influenced by Federal decisions. However, industry planners are likely to discount agency plans that project specific long term needs for Class VIII systems as being far too uncertain to provide a basis for decision.

The options for government stimulation of this area range from maintaining the vendors' confidence that a Federal market will continue in the next decade, to direct support of the necessary R and D, to issuing an rfp for Class VIII systems. The members of the panel did not appear to support the last two options. There may be other immediate options that fit the role of other agencies; however, NSF can help maintain a strong U.S. position in the field by supporting a strong program of:

- basic research in universities in advanced computing concepts, including the software, algorithms, architecture, subsystems and components.
- basic research in problems that require major computer systems;
- the training of scientists and engineers who will be familiar with the concepts, design and use of major computer systems.
- assured access for researchers and graduate students to major computer systems.

The NSF is likely to have a direct, although minor, influence on the market through supporting the use of Class VIII systems. The basic research supported by NSF underpins the development of systems beyond Class VIII.

- 3 -

Access to Existing Class VI Computers

The computing needs for university research include access to supercomputers, specialized computational equipment and array processors with supermini-computers. This memorandum considers the rising need for university researchers in a variety of disciplines to have access to up-to-date large computers. That such a need exists is the opinion of NSF staff involved with research areas requiring computation and is supported by the Lax report and other studies specific to certain disciplines. Opportunities already exist for researchers in some disciplines to access Class VI machines, for example, through NCAR and the MFE network. Additionally, several universities and some industrial concerns have either acquired or plan to acquire CRAY I or Cyber 205 machines and will sell time. Federal laboratories also have time available for outside use.

NSF recommends that an interagency working group be established, under the FCCSET, to examine options for meeting an aggregated demand for access to existing Class VI machines. The group will need inputs from experts in industry and the universities. The options include:

- providing access to machines in Federal laboratories;
- providing access to existing university based Class VI machines;
- the purchase of commercial computing services;
- the establishment of a dedicated network or center.

NSF recognizes that it will have a significant role in this planning process and in the implementation of recommended approaches to solving the problem.

Coordination

Requirements for coordination through more permanent working groups, or at the policy level, will not be clear until the FCCSET supercomputer panel has reached firm conclusions on the first two issues addressed in your memo.

However, in addition to a group charged with examining the access question, NSF believes that coordination of some specific elements of computer R and D, e.g., computer architecture may be important. NSF does not perceive a need to strengthen coordination of broad research programs in computer science and engineering which is already active and effective on an informal basis.